

Detecting Irregular Curvilinear Structures in Gray Scale and Color Imagery using Multi-Directional Oriented Flux

Engin Türetken* Carlos Becker Przemysław Głowacki Fethallah Benmansour Pascal Fua
CVLab, EPFL, Lausanne, Switzerland

{engin.turetken, carlos.becker, przemyslaw.glowacki, fethallah.benmansour, pascal.fua}@epfl.ch

Abstract

We propose a new approach to detecting irregular curvilinear structures in noisy image stacks. In contrast to earlier approaches that rely on circular models of the cross-sections, ours allows for the arbitrarily-shaped ones that are prevalent in biological imagery. This is achieved by maximizing the image gradient flux along multiple directions and radii, instead of only two with a unique radius as is usually done. This yields a more complex optimization problem for which we propose a computationally efficient solution.

We demonstrate the effectiveness of our approach on a wide range of challenging gray scale and color datasets and show that it outperforms existing techniques, especially on very irregular structures.

1. Introduction

Enhancing curvilinear structures is a crucial preprocessing step in a broad array of biomedical delineation tasks, ranging from finding blood vessels [4, 15, 3] to dendritic arbors [23, 24, 21]. The most popular techniques rely on filters that are crafted to respond to locally linear features [13, 7, 20, 18], optimized for specific profiles [11, 22], or learnt [19, 10]. They return a measure that, ideally, should be maximal on the centerline of tubular structures.

In recent years, Optimally Oriented Flux (OOF) [13] has emerged as one of the best such *tubularity measures*, in part because it can easily be normalized for scale and is relatively insensitive to the presence of adjacent structures. Thus, it provides a higher discriminative power in detecting curvilinear structures in scale-space than Hessian based features. However, it also produces a strong response on high contrast edges and the Oriented Flux Antisymmetry (OFA) measure [14] was later introduced to address this issue. Combining OOF and OFA has been shown to be very

*This paper was supported in part by the EU ERC project MicroNano and in part by the Swiss National Science Foundation.




Figure 1. Sources of irregularity in biological imagery, best viewed in color. Each one of the four examples shown here includes at the top a maximal or minimal intensity projection and below two cross sections of tubular structures. These cross sections correspond to the yellow lines overlaid on the top image. (a) Confocal stack of blood vessels with non-uniform staining. (b) Brainbow stack of neurites [17] with severe occlusion and image noise. (c,d) Bright-field image stacks with irregular staining and point-spread function blur.

effective on relatively regular tube-like structures.

However, performance is adversely affected when cross sections become very irregular and deviate strongly from being circular. As shown in Fig. 1, this is an issue because curvilinear structures in biological imagery often appear as jagged filaments whose cross sections are extremely irregular, due among others to imaging noise, artifacts, non-uniform staining, occlusion and point-spread function.

In this paper, we introduce a new tubularity measure that we refer to as Multi-Directional Oriented Flux (MDOF),

designed to handle this kind of irregularities. We demonstrate that it is more effective at handling them not only than the OOF and OFA combination but also than other popular methods [16, 7].

This is achieved by maximizing the image gradient flux along multiple directions and radii, rather than only along two orthogonal directions with the same radii as in the OOF framework. This represents a more complicated optimization problem for which we propose a computationally efficient solution. For color images, we further boost performance by exploiting, at each image location, gradient information of similar colors from its vicinity.

Our contribution is both this new MDOF measure for detecting curvilinear structures in gray scale and color images and an approach to computing it efficiently, which makes it practical to solve real-world problems. We will demonstrate its effectiveness on four difficult and very different datasets.

2. Related Work

Hand-designed curvilinear structure detectors fall into two main categories: the very popular Hessian-based detectors and model-based optimal filters. Here, we describe them and discuss their advantages and drawbacks. Another significant group of techniques are those that learn filters [19, 10] given ground-truth data. In practice, acquiring such data is an expensive and laborious process and we focus here on techniques that do not require it.

2.1. Hessian-Based Detectors

Many methods rely on computing second order derivatives to detect filament-like patterns. The main assumption is that the desired structure is characterized by a local principal direction and a cross-sectional plane with Gaussian intensity profile. The most common tool to characterize such patterns is the Hessian matrix. Hessian-based approaches [7, 20, 6, 18] consist in computing the eigenvalues of the Hessian and combining them to obtain a tubularity measure that quantifies the likelihood that a specific pixel lies on a centerline of a filament-like structure. Despite their popularity, Hessian-based approaches suffer from three major drawbacks:

1. Sensitivity to local deformations. If the cross-section does not present a circular profile the assumptions justifying using the Hessian matrix no longer hold, which often leads to false negative responses. The lack of robustness in the presence of an irregular profile is a general drawback common to all existing hand designed detectors.
2. Detection of thick branches requires a large variance of the Gaussian, which induces substantial blur. Therefore the response of such filters at a thick branch may

be biased by the presence of adjacent structures in the immediate neighborhood.

3. Last but not least, normalization over scales of Hessian-based tubularity measures is not straightforward. A normalization procedure for 2D Gaussian ridge profiles was introduced in [16]. Unfortunately, curvilinear structures in real data do not necessarily conform to such an appearance model. Moreover, for datasets where the cross-sections do not have Gaussian profiles the normalization has a tendency to attenuate drastically the responses at large scales. Consequently, multi-scale Hessian-based measures favor small scales [5] and do not provide reliable scale estimates.

2.2. Optimal Filtering

Jacob et al. [11, 1] propose a method utilizing a class of steerable filters [8]. The proposed set of filters are Gaussian derivatives of up to, typically, the 4th order. A related approach, the Optimally Oriented Flux filtering [13], consists of convolving the second order derivatives of the image with the indicator of a sphere, which is a steerable filter designed for detecting curvilinear structures. Compared to Hessian-based detectors, the Oriented Flux filter is straightforward to normalize over scale. The OOF computation is confined to a local spherical region which results in less sensitivity to the presence of adjacent structures as compared to Hessian-based methods that involve convolution with the second derivative of Gaussian. This approach has been extended to take into account antisymmetry of the projected gradients at the boundaries of tubular structures [14], as explained in detail in Section 3.1.

All these approaches assume an ideal circular appearance of the curvilinear structure cross-sections, which is often not true in biomedical images, as shown in Fig. 1. More sophisticated model-based approaches such as the super-ellipsoid fitting method [22], provide better robustness to these irregularities. However, they are not particularly suitable for dense computation since their filters cannot be steered and they are hard to optimize.

3. Approach

Our objective is to devise an algorithm able both to reliably detect the centerlines of irregularly shaped curvilinear structures and to estimate their radius. To this end, we assign to each spatial location and each radius value a *tubularity score* and then extract the centerline points and their radii as local maxima in scale-space. Formally, we compute a tubularity value $f(\mathbf{x}, r)$ for each 3D image location \mathbf{x} and radius r . It quantifies the likelihood of \mathbf{x} being on the centerline of a curvilinear structure of radius r within a range $[r_{min} r_{max}]$. In practice, we sample the radius range




Figure 2. Representative results on four image stacks, each corresponding to one of the modalities described in Section 4.1. From top to bottom, they are Retina, Brainbow, Brightfield, and VC6. (a) Original image stacks. (b) Maximum intensity projection of the 4-D scale-space tubularity volumes obtained with our approach along the scale dimension. The images are color-coded so that higher scale values appear more saturated, going from white to red. (c) Maximum intensity projection of the non-maxima suppressed tubularity volumes obtained by running the NMST algorithm of Section 4.2. (d) Segmentation results obtained by maximizing the Jaccard index of Section 4.2. The 3D renderings are created using the publicly available Vaa3D software.

uniformly by increments of one pixel spacing. We therefore produce 4D tubularity images, such as those depicted by Fig. 2(b). Without loss of generality, from now on, we will assume that curvilinear structures are brighter than the background. If they are darker, the image gray levels can simply be reversed. Furthermore, we will use the terms radius and scale interchangeably.

Our tubularity measure extends the well-known Optimally Oriented Flux (OOF) [13] filter, which we first review briefly for completeness. We then introduce our approach to handling the irregularity of many curvilinear structures, which deviates from the idealized models used by the OOF techniques. Finally, we extend our approach to color images.

3.1. Optimally Oriented Flux

As discussed in Section 2, many existing approaches to enhancing curvilinear structures [16, 7, 20] rely on a Hessian based measure obtained by convolving the image with second order Gaussian derivatives. This involves some amount of smoothing. As a result, they take into account image intensities in the vicinity of the structures, which can adversely affect their accuracy in the presence of adjacent structures. Furthermore, in order to estimate the structure radius from the Gaussian standard deviation, they use an idealized intensity profile model, which is not trivial to obtain and is not applicable to all imaging modalities.

The OOF filter addresses these issues by considering intensity values only within a spherical volume of certain diameter, which provides a good estimate of the structure width [13]. The filter is computed by convolving the second derivatives of the image with the indicator function of the sphere. More formally, its value $f(\mathbf{x}, \mathbf{p}, r)$ for image location \mathbf{x} , radius r , and orientation \mathbf{p} is obtained by integrating the projected image gradients in the \mathbf{p} direction over a sphere ∂S_r of radius r centered at \mathbf{x} . This is written as

$$f(\mathbf{x}, \mathbf{p}, r) = \frac{1}{4\pi r^2} \int_{\partial S_r} ((\nabla(G_{\sigma_0} * I)(\mathbf{x} + \mathbf{h}) \cdot \mathbf{p})\mathbf{p}) \cdot \mathbf{n} da, \quad (1)$$

where \mathbf{n} is the outward unit normal of ∂S_r , G_{σ_0} is a regularizing Gaussian with a small standard deviation σ_0 typically equal to voxel spacing, $\mathbf{h} = r\mathbf{n}$ is the position vector with its tip on ∂S_r , and da is the infinitesimal area on ∂S_r . The smaller the value of $f(\mathbf{x}, \mathbf{p}, r)$, the more likely it is that \mathbf{x} is the center of a tube with radius r and orientation \mathbf{p} .

For each \mathbf{x} , we therefore look for the values of \mathbf{p} and r that minimize f . It can be shown that $f(\mathbf{x}, \mathbf{p}, r)$ can be rewritten as the quadratic form $\mathbf{p}^T Q_{\mathbf{x},r} \mathbf{p}$, where $Q_{\mathbf{x},r}$ is known as the *oriented flux matrix* [13]. Using the divergence theorem, its entries can be expressed as

$$Q_{\mathbf{x},r}^{i,j} = \frac{1}{4\pi r^2} (\partial_{i,j} G_{\sigma_0}(\mathbf{x}) * \mathbb{1}_r * I(\mathbf{x})), \quad (2)$$

where $\mathbb{1}_r$ is the indicator of the sphere of radius r . The OOF tubularity measure is defined as the sum of the two OOF filter responses evaluated along two orthogonal directions \mathbf{p}_1 and \mathbf{p}_2 defining the structure's cross-sectional plane [13]:

$$\begin{aligned} f_{OOF}(\mathbf{x}, r) &= \max_{\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_1 \perp \mathbf{p}_2} -f(\mathbf{x}, \mathbf{p}_1, r) - f(\mathbf{x}, \mathbf{p}_2, r) \\ &= \max_{\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_1 \perp \mathbf{p}_2} -\mathbf{p}_1^T Q_{\mathbf{x},r} \mathbf{p}_1 - \mathbf{p}_2^T Q_{\mathbf{x},r} \mathbf{p}_2, \quad (3) \end{aligned}$$

which can be shown to be equal to the negative sum of the two smallest eigenvalues of $Q_{\mathbf{x},r}$. Their associated eigenvectors $\mathbf{e}_{\mathbf{x},r}^1$ and $\mathbf{e}_{\mathbf{x},r}^2$ provide the optimal directions for the above maximization.

Due to its intrinsic symmetry, the OOF filter yields high responses along centerlines of curvilinear structures at their associated scales. However, it also responds strongly to edges. To prevent this, a gradient antisymmetry function $g(\mathbf{x}, \mathbf{p}, r)$ was introduced in [14]:

$$g(\mathbf{x}, \mathbf{p}, r) = \frac{1}{4\pi r^2} \int_{\partial S_r} (\nabla(G_{\sigma_0} * I)(\mathbf{x} + \mathbf{h}) \cdot \mathbf{p}) da, \quad (4)$$

which is equal to the inner product $\mathbf{p}^T \mathbf{q}_{\mathbf{x},r}$. The term $\mathbf{q}_{\mathbf{x},r}$ is called the *oriented flux antisymmetry vector* (OFA) whose entries are computed as follows:

$$\mathbf{q}_{\mathbf{x},r}^i = \frac{1}{4\pi r^2} (\partial_i G_{\sigma_0}(\mathbf{x}) * \mathbb{1}_r * I(\mathbf{x})), \quad (5)$$

The absolute value of the antisymmetry function takes high values at structure boundaries and vanishes on perfectly symmetric shapes, for instance, along the centerline of a tube. Using this principle, the OOF tubularity measure can be refined to take into account structure symmetry as follows [14]:

$$\begin{aligned} f_{OFA}(\mathbf{x}, r) &= f_{OOF}(\mathbf{x}, r) - \sqrt{g(\mathbf{x}, \mathbf{e}_{\mathbf{x},r}^1, r)^2 + g(\mathbf{x}, \mathbf{e}_{\mathbf{x},r}^2, r)^2} \\ &= f_{OOF}(\mathbf{x}, r) - \sqrt{\mathbf{e}_{\mathbf{x},r}^{1T} H_{\mathbf{x},r} \mathbf{e}_{\mathbf{x},r}^1 + \mathbf{e}_{\mathbf{x},r}^{2T} H_{\mathbf{x},r} \mathbf{e}_{\mathbf{x},r}^2}, \quad (6) \end{aligned}$$

with $H_{\mathbf{x},r} = \mathbf{q}_{\mathbf{x},r} \mathbf{q}_{\mathbf{x},r}^T$. The two antisymmetry terms above act complementary to the OOF filter as they annihilate the OOF response away from the structure centerlines. However, since both OOF and OFA are evaluated only at two directions and a single radius value, the measure of Eq. 6 tends to favor circular and symmetric cross sections.

3.2. MDOF: Multi-Directional Oriented Flux

Curvilinear structures such as the ones shown in Fig. 1, often appear as irregular filaments in medical imagery due to reasons such as imaging noise, non-uniform staining and point-spread function. We address this issue by finding multiple directions and radii of maximal response of the joint

oriented flux and antisymmetry operators. More specifically, given a set of predetermined radius levels S_r , we compute a number of direction and radius pairs ($\mathbf{p}_i \in \mathbb{R}^3, r_i \in S_r$) that maximize

$$\begin{aligned} f_{MS}(\mathbf{x}) &= - \sum_{i=1}^d \mathbf{p}_i^T (Q_{\mathbf{x},r_i} + \alpha H_{\mathbf{x},r_i}) \mathbf{p}_i \quad (7) \\ &= - \sum_{i=1}^d \mathbf{p}_i^T F_{\mathbf{x},r_i} \mathbf{p}_i, \end{aligned}$$

subject to the constraint that all \mathbf{p}_i lie on the same plane at regular angular intervals.




Figure 3. MDOF takes into account non-circular cross-sections by allowing multiple directions \mathbf{p}_i and radii r_i at every possible location \mathbf{x} .

A schematic of such distribution is presented in Fig. 3, where the two first orientations and radii are shown in dark red, and additional ones in dashed red.

As in Eq. 6, the matrix $H_{\mathbf{x},r_i}$ is equal to the quadratic form $\mathbf{q}_{\mathbf{x},r_i} \mathbf{q}_{\mathbf{x},r_i}^T$ of the antisymmetry vector as defined in Eq. 5 and α is a regularization parameter that determines the amount of symmetry enforced along each direction. It is set to 0.1 in all our experiments.

Note that, in contrast to [14]

which first finds the two dominant eigenvalues of $Q_{\mathbf{x},r}$ and then plugs their eigenvectors into the antisymmetry term, we jointly optimize for the oriented flux and the antisymmetry terms.

An important consideration when solving Eq. 7 is to do it fast because it needs to be done for each individual voxel. To this end, we employ a greedy approach, which first finds the dominant direction \mathbf{p}_1^* and the associated radius r_1^* that maximizes $f_{MS}(\mathbf{x})$ for $d = 1$ and then (\mathbf{p}_2^*, r_2^*) that maximizes the same function subject to the constraint that \mathbf{p}_2^* is perpendicular to \mathbf{p}_1^* . The first step involves computing \mathbf{p}_1^* as the eigenvector associated with the smallest eigenvalue among those of $F_{\mathbf{x},r_i}$, for all $r_i \in S_r$. Although solving the second step requires satisfying the linear constraints $\mathbf{p}_2^T \mathbf{p}_1^* = 0$ and $\mathbf{p}_2^T \mathbf{p}_2 = 1$, it can be done in closed form by solving the following eigenvalue problem [9]:

$$(\mathbf{e}^*, r_2^*) = \operatorname{argmax}_{\mathbf{e}, r_i \in S_r} - \mathbf{e}^T P F_{\mathbf{x},r_i} P \mathbf{e}, \quad (8)$$

where the matrix P is defined as $P = I - \mathbf{p}_1^* \mathbf{p}_1^{*T}$. The vector \mathbf{e}^* can be computed as the eigenvector associated with the smallest eigenvalue among those of $P F_{\mathbf{x},r_i} P$, for all $r_i \in S_r$. The optimal second direction \mathbf{p}_2^* can then be found as $\mathbf{p}_2^* = P \mathbf{e}^*$ [9].

We then find the additional $d - 2$ directions by sampling them at equiangular intervals to span π radians in the plane defined by \mathbf{p}_1^* and \mathbf{p}_2^* . The corresponding radius values are found by maximizing the same cost function of Eq. 7 in the set S_r . This produces a 3-D score image, which we then extend to 4-D by adding the strongest response of $\mathbf{p}_i^T F_{\mathbf{x},r_i} \mathbf{p}_i$ for each radius level. We take the final score for radius r at location \mathbf{x} to be

$$f_{MDOF}(\mathbf{x}, r) = -\lambda_{\mathbf{x},r}^1 + \frac{1}{d} \max_{\{\mathbf{p}_i, r_i\}_{i=1}^d} f_{MS}(\mathbf{x}), \quad (9)$$

where $\lambda_{\mathbf{x},r}^1$ is the smallest eigenvalue of $F_{\mathbf{x},r}$. At each radius level, we use only a single eigenvalue, which gives the strongest response, because, in most datasets, the gradient information is strong and provides a reliable estimate only along one direction in the irregular cross-sectional profiles. For instance, in the brightfield stacks of Fig. 1(c,d), this direction usually lies in the lateral (x-y) plane of the in-focus z slice. Taking into account the gradient in out-of-focus regions therefore usually results in overestimation of the scale and mislocalization of the structure centerline.

The second term in Eq. 9 measures the likelihood of \mathbf{x} being on the centerline of a curvilinear structure, while the first one measures the contribution of the radius level r . As a result, at each image point \mathbf{x} , the radius level r that yields the strongest gradient flux on the associated sphere S_r will be assigned the highest score. As we will show in Section 4, this measure is more effective in suppressing spurious responses away from the centerlines when the structures are irregular.

3.3. Handling Color Images

With the advent of new imaging techniques, color information has gained increasing importance in the analysis of complex curvilinear structures. The Brainbow [17] and STORM [12] techniques, for instance, allow labelling neuron cells in a tissue slice with multiple colors using a random mixture of few fluorescent proteins. This results in each neuron being labelled with a sufficiently distinct color, which is nearly constant along its axon and dendrites. In fact, the color constancy property also widely holds for other curvilinear structures such as road networks.

We use this property to compute, for each image location \mathbf{x} , a similarity image, which signifies how likely both \mathbf{x} and another location \mathbf{y} in its vicinity belong to the same curvilinear structure based on their color. We then use the resulting similarity image to compute the tubularity measure of Eq. 9 at location \mathbf{x} . More precisely, given an RGB image stack I , we first convert it into the perceptually uniform CIELAB space. This is done because the aforementioned color imaging techniques aim to maximize the color contrast among curvilinear structures for visual inspection by human experts. We then compute the similarity image J_x

for \mathbf{x} as a function $s(\cdot)$ of the color distance between \mathbf{x} and \mathbf{y} , weighted by the value of the brightness image L , which is equivalent to $\max(R,G,B)$. We write

$$J_{\mathbf{x}}(\mathbf{y}) = L(\mathbf{y}) s(I(\mathbf{x}), I(\mathbf{y})), \quad \forall \mathbf{x}. \quad (10)$$

In this work, we take the similarity function s to be a multivariate Gaussian distribution with mean vector $I(\mathbf{x})$ and a diagonal covariance matrix with diagonal entries σ . In practice, computing the similarity image for each location is quite expensive. That is why we limit the number of similarity images to a number of color clusters, which we compute using the k-means algorithm. Let C be the set of resulting cluster centroids. Then, Eq. 10 becomes:

$$J_{\mathbf{x}}(\mathbf{y}) = L(\mathbf{y}) s(c_{h(\mathbf{x})}, I(\mathbf{y})), \quad \forall c_{h(\mathbf{x})} \in C, \quad (11)$$

where $h(\mathbf{x})$ is a function that returns the cluster label for location \mathbf{x} . Note that, Eq. 11 requires sequentially computing only $|C|$ similarity images, which serve as input to the tubularity computation described in the previous section. The result is that, for each location, only image evidence from nearby locations of similar color are taken into account.

4. Results

In this section, we first describe the datasets we used for evaluation purpose. We then introduce our evaluation metrics and present our results.

4.1. Datasets and Parameters

We evaluated our approach on the four different datasets described below. In our experiments, the radius levels are measured in the units of image spacing along the XY axes.

- **Retina** Two confocal micrographs of direction selective retinal ganglion cells and loopy vasculature networks. In our computations, we sampled 9 radius values ranging from 3 to 16.
- **Brainbow** Four micrographs of mice primary visual cortex acquired using the brainbow staining technique [17]. We sampled 9 radius values ranging from 3 to 11.
- **Brightfield** Four brightfield micrographs of biocytin-stained rat brains. We sampled 12 radius values ranging from 1 to 12.
- **VC6** Four brightfield micrographs of biocytin-labeled cat primary visual cortex layer 6 taken from the DIADEM challenge data [2]. We sampled 12 radius values ranging from 1 to 12.

We used the same parameters mentioned in Section 3.2 for all the four datasets: the number of direction and radius pairs $d = 10$, color similarity sigma $\sigma = 50$ in the the CIELAB space and the number of color clusters $|C| = 50$.

4.2. Evaluation Metrics

We compare our approach against Optimally Oriented Flux (OOF) [13], Oriented Flux with Oriented Flux Anti-symmetry (OOF/OFA) [14], and the older Hessian-Based Frangi [7] and Lindeberg [16] approaches.

To this end we begin by computing the 4-D tubularity image stacks, where the fourth dimension represents radius, using all five methods. These stacks can be turned into ordinary 3-D stacks by choosing the radius that yields the highest tubularity value at each location. We then perform non maxima suppression on the resulting 3-D stacks so that we can evaluate them both in terms of how close the maxima are from ground-truth centerlines and in terms of voxel-wise accuracy when segmenting the tubular structures by jointly using estimated centerline location and radius.

Non Maxima Suppression. All five methods compute a matrix whose eigenvectors can be used to compute a local orientation, which we use in conjunction with the radii estimates to implement a Canny-style approach to non-maxima suppression. More specifically, we suppress voxels that are not local maxima in the plane that is perpendicular to the local orientation and within the circular neighborhood defined by their radius.

Unfortunately, because the estimated orientations are not always accurate, this often removes some of the legitimate centerline points and produces gaps along the structures, which would skew our accuracy estimates. We solve this problem by first finding the Minimum Spanning Tree (MST) that connects all the voxels in the original 3-D tubularity stack. We use 26-connectivity and take the weight of each edge to be the negative exponential of the mean tubularity score of the voxels it connects. We then find connected components in the non-maxima suppressed stack and identify the MST paths that link pairs of these. For each component, we retain only the two MST paths that have the highest average tubularity scores. This yields *Non-Maxima Suppressed Trees* (NMST) such as the ones depicted by the third column of Fig. 2.

Centerline Detection Accuracy. To perform precision-recall (PR) analysis, given an NMST and a ground-truth centerline, we introduce a tolerance factor τ measured in XY spacing units, such that the centerlines present in the NMST are considered true positives when they are at most τ distance away from a ground truth centerline voxel. The resulting PR curves are plotted for $\tau = 2$ in Fig. 4(a) and additional ones are provided as supplementary material to show that their ordering is relatively insensitive to τ .

Joint Centerline and Radius Accuracy. Note that locally maximum voxels in the NMST are still attached tubularity values. To also estimate the quality of our radius estimates, we therefore threshold it at different values. For

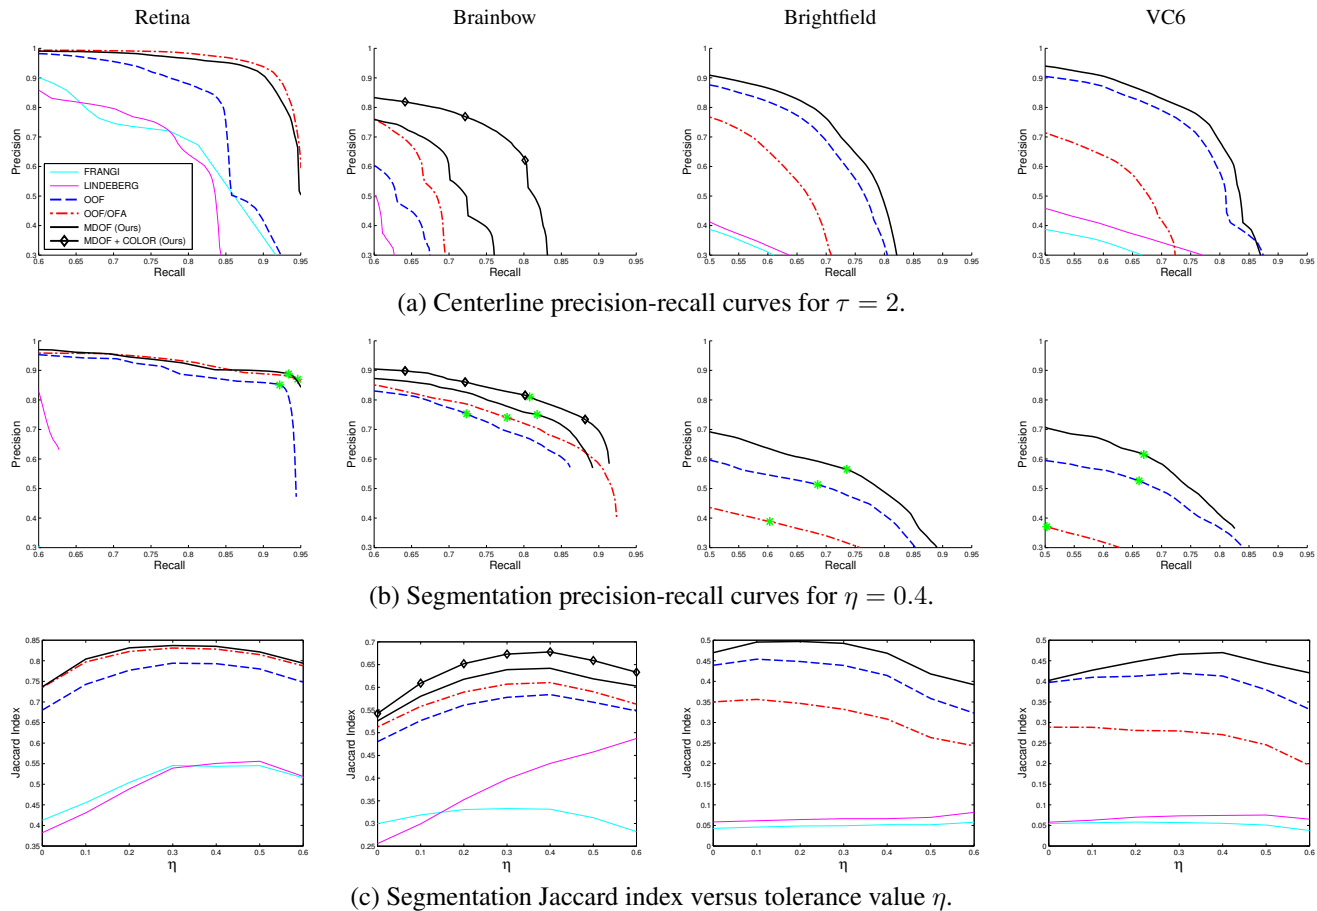


Figure 4. Performance curves for our approach and the baselines on four different datasets. The green stars in the middle row indicate where the Jaccard index is maximized.




Figure 5. Tubular structure segmentation from centerline and scale estimates. Spheres are centered at NMST image voxels that above an established threshold, generating the segmentation on the right.

each, we then construct a tubular structure by using our radius estimates to place spheres at each non-empty voxel of the thresholded NMST as shown in Fig. 5. This produces a full segmentation that we compare voxel-wise to the ground-truth. To handle potential inaccuracies in the ground-truth, which are always present, we define an exclusion zone $[(1 - \eta)r, (1 + \eta)r]$ with $\eta < 1$ and ignore voxels that are closer than ηr from the surface of the ground truth tube, where r is its radius. We set $0 \leq \eta \leq 0.6$ and obtain different PR curves for each η value. In Fig. 4(b), we plot those we obtain for $\eta = 0.4$. Finally, in Fig. 4(c), we

	LINDBERG	FRANGI	OOF	OOF/OFA	MDOF	MDOF COLOR
Retina	101	95	240	393	1344	N/A
Brainbow	10	9	18	30	107	634
Brightfield	132	124	236	384	1576	N/A
VC6	45	42	84	136	553	N/A

Table 1. Average run-times per image stack measured in seconds on a multi-core machine.

show the best Jaccard index as a function of η . Note that the Frangi and Lindeberg measures are not visible in some figures because they achieve low performance values outside the presented range.

4.3. Interpretation

In all our datasets, the three top contenders clearly are our approach, OOF, and OOF/OFA. In the Retina stacks, the tubular structures are relatively regular with only few vascular branches containing non-uniform staining artifacts. Unsurprisingly, all three perform similarly, with OOF/OFA doing best in terms of centerline detection accuracy and MDOF doing slightly better in terms of joint centerline and radius in the operating range at high recall rates. Note that

they both perform better than OOF alone, thus demonstrating the importance of the antisymmetry term.

In the three other datasets the tubular structures are much more irregular and MDOF clearly outperforms both OOF/OFA and OOF. Note also that in the Brightfield and VC6 case, unlike in the Retina case, OOF performs considerably better than OOF/OFA, which indicates that the antisymmetry term loses some of its power when the structures become highly irregular. In other words, in all the three datasets, MDOF performs either better or much better than either OOF and OOF/OFA taken individually.

Finally, for the Brainbow dataset, which is the only one in which the stacks are in color, using the color information yields a further performance boost.

These improvements in performance, however, come at the expense of higher running times as listed in Table. 1. Although OOF, OOF/OFA and MDOF have the same computational complexity with respect to the image size, the joint search for the radii and directions in MDOF incurs a significant overhead.

Our method and the baselines are implemented using the ITK library and will be published as open-source software.

5. Conclusion

We presented and validated a new tubularity measure that performs better than existing approaches on irregular structures whose cross sections deviate from circular symmetric profiles. This is important because many imaging modalities produce irregular structures as a result of noise, due to point spread function blur, and non-uniform staining, among others. Furthermore, we showed that our approach can be extended to exploit color information when available to increase robustness.

References

- [1] F. Aguet, M. Jacob, and M. Unser. Three-Dimensional Feature Detection Using Optimal Steerable Filters. In *ICIP*, September 2005.
- [2] G. A. Ascoli, K. Svoboda, and Y. Liu. Digital Reconstruction of Axonal and Dendritic Morphology Diadem Challenge, 2010. <http://diademchallenge.org/>.
- [3] F. Benmansour and L. Cohen. Tubular Structure Segmentation Based on Minimal Path Method and Anisotropic Enhancement. *IJCV*, 92(2):192–210, 2011.
- [4] L. D. Cohen and T. Deschamps. Grouping Connected Components Using Minimal Path Techniques. Application to Reconstruction of Vessels in 2D and 3D Images. In *CVPR*, pages 102–109, 2001.
- [5] O. Dzyubak and E. Ritman. Automation of Hessian-Based Tubularity Measure Response Function in 3D Biomedical Images. *International Journal of Biomedical Imaging*, 2011.
- [6] A. H. Foruzan, R. A. Zoroofi, Y. Sato, and M. Hori. A Hessian-based filter for vascular segmentation of noisy hepatic CT scans. *IJCARS*, 7(2):199–205, 2012.
- [7] A. Frangi, W. Niessen, K. Vincken, and M. Viergever. Multiscale Vessel Enhancement Filtering. *Lecture Notes in Computer Science*, 1496:130–137, 1998.
- [8] W. Freeman and E. Adelson. The Design and Use of Steerable Filters. *PAMI*, 13:891–906, 1991.
- [9] G. Golub. Some Modified Matrix Eigenvalue Problems. *SIAM Review*, 15(2):318–334, 1973.
- [10] G. Gonzalez, F. Aguet, F. Fleuret, M. Unser, and P. Fua. Steerable Features for Statistical 3D Dendrite Detection. In *MICCAI*, pages 625–32, September 2009.
- [11] M. Jacob and M. Unser. Design of Steerable Filters for Feature Detection Using Canny-Like Criteria. *PAMI*, 26(8):1007–1019, August 2004.
- [12] M. Lakadamyali, H. Babcock, M. Bates, X. Zhuang, and J. Lichtman. 3D Multicolor Super-Resolution Imaging Offers Improved Accuracy in Neuron Tracing. *PLoS ONE*, 7(1):e30826, 2012.
- [13] M. Law and A. Chung. Three Dimensional Curvilinear Structure Detection Using Optimally Oriented Flux. In *ECCV*, 2008.
- [14] M. Law and A. Chung. An Oriented Flux Symmetry Based Active Contour Model for Three Dimensional Vessel Segmentation. In *ECCV*, pages 720–734, 2010.
- [15] H. Li and A. Yezzi. Vessels as 4-D Curves: Global Minimal 4D Paths to Extract 3-D Tubular Surfaces and Centerlines. *TMI*, 26(9):1213–1223, 2007.
- [16] T. Lindeberg. Edge Detection and Ridge Detection with Automatic Scale Selection. In *CVPR*, page 465, 1996.
- [17] J. Livet, T. Weissman, H. Kang, R. Draft, J. Lu, R. Bennis, J. Sanes, and J. Lichtman. Transgenic strategies for combinatorial expression of fluorescent proteins in the nervous system. *Nature*, 450(7166):56–62, 2007.
- [18] E. Meijering, M. Jacob, J.-C. F. Sarría, P. Steiner, H. Hirling, and M. Unser. Design and Validation of a Tool for Neurite Tracing and Analysis in Fluorescence Microscopy Images. *Cytometry Part A*, 58A(2):167–176, April 2004.
- [19] A. Santamaría-Pang, T. Bildea, C. M. Colbert, P. Saggau, and I. Kakadiaris. Towards Segmentation of Irregular Tubular Structures in 3D Confocal Microscope Images. In *MICCAI MIAAB*, 2006.
- [20] Y. Sato, S. Nakajima, H. Atsumi, T. Koller, G. Gerig, S. Yoshida, and R. Kikinis. 3D Multi-Scale Line Filter for Segmentation and Visualization of Curvilinear Structures in Medical Images. *MIA*, 2:143–168, June 1998.
- [21] E. Türetken, F. Benmansour, and P. Fua. Automated Reconstruction of Tree Structures Using Path Classifiers and Mixed Integer Programming. In *CVPR*, June 2012.
- [22] J. Tyrrell, E. di Tomaso, D. Fuja, R. Tong, K. Kozak, R. Jain, and B. Roysam. Robust 3D Modeling of Vasculature Imagery Using Superellipsoids. *MIA*, 26(2):223–237, 2007.
- [23] X. Wang, B. Wang, and L. Zhang. Airport Detection in Remote Sensing Images Based on Visual Attention. In *ICONIP*, 2011.
- [24] T. Zhao, J. Xie, F. Amat, N. Clack, P. Ahammad, H. Peng, F. Long, and E. Myers. Automated Reconstruction of Neuronal Morphology Based on Local Geometrical and Global Structural Models. *Neuroinformatics*, 9:247–261, May 2011.